

Claims

What is claimed is:

1. A system comprising:
 - a cluster having a plurality of nodes wherein at least one of the nodes is a candidate node;
 - a plurality of resource groups;
 - a clustering mechanism executing on the cluster configured to activate a first resource group of the plurality of resource groups on the candidate node;
 - and
 - a resource group affinity of the plurality of resource groups, wherein the resource group affinity comprises a unidirectional association between the first resource group of the plurality of resource groups and a second resource group of the plurality of resource groups.
2. The system of claim 1, wherein the plurality of resource groups comprises:
 - a plurality of resources configured to provide at least one service; and
 - a monitor configured to observe the activity of the at least one service.
3. The system of claim 1, wherein the resource group affinity further comprises:
 - an affinity type; and
 - an affinity strength.
4. The system of claim 3, wherein the affinity strength is one selected from the group consisting of strong, weak, and mutual strong and the affinity type is one selected from the group consisting of positive and negative.
5. The system of claim 3, wherein the mutual strong affinity strength comprises only a positive affinity type.

6. The system of claim 3, wherein the first resource group is activated on a same node as the second resource group when the resource group affinity comprises one selected from a strong and mutual strong affinity strength and a positive affinity type.
7. The system of claim 6, wherein the first resource group follows the second resource group if the second resource group changes nodes or goes offline.
8. The system of claim 3, wherein the first resource group is activated on a different node than the second resource group when the resource group affinity comprises a strong affinity strength and a negative affinity type.
9. The system of claim 8, wherein the first resource group relocates to another node if the second resource group moves to the same node as the first resource group.
10. The system of claim 3, wherein the first resource group may be activated on a same node as the second resource group when the resource group affinity comprises a weak affinity strength and a positive affinity type, but if an attempt fails, the first resource group may be activated on a different node than the second resource group.
11. The system of claim 3, wherein the first resource group may be activated on a different node than the second resource group when the resource group affinity comprises a weak affinity strength and a negative affinity type, but if an attempt fails, the first resource group may be activated on a same node as the second resource group.
12. The system of claim 3, wherein the first resource group is activated on a same node as the second resource group and the first resource group is able to move the second resource group to a different node when the resource group affinity comprises a mutual strong affinity strength and a positive affinity type.

13. The system of claim 1, wherein one of the plurality of resource groups is activated prior to any other resource group if the one of the plurality of resource groups has no resource group affinity.
14. The system of claim 1, wherein a least constrained resource group of the plurality of resources groups is activated prior to more constrained resource groups.
15. The system of claim 14, wherein a first of the plurality of resource groups has no strong affinities for any other resource group and a second of the plurality of resource groups has no strong affinities for any resource group that follows the second of the plurality of resource groups.
16. A method to activate a first resource group comprising:
 - determining a plurality of candidate nodes for the first resource group;
 - eliminating any of the plurality of candidate nodes that violates a strong affinity of the first resource group, to generate a plurality of affinity candidate nodes;
 - determining a composite value for each of the plurality of affinity candidate nodes;
 - and
 - sorting the plurality of affinity candidate nodes based on a pre-defined priority using the composite values to obtain a sorted list.
17. The method of claim 15, further comprising:
 - activating the first resource group on at least one of the plurality of affinity candidate nodes based on the sorted list.
18. The method of claim 15, wherein determining the composite value further comprises:
 - determining a number of resource groups that are removed from one of the plurality of affinity candidate nodes when the first resource group is activated on the one of the plurality of candidate nodes;

calculating the net number of outgoing weak affinities satisfied when the first resource group is activated on the one of the plurality of candidate nodes;
calculating the net number of incoming weak affinities satisfied when the first resource group is activated on the one of the plurality of candidate nodes;
and
determining an ordinal position of the one of the plurality of candidate nodes in the first resource group's nodelist property.

19. The method of claim 18, wherein the resource groups that are removed comprise a strong affinity for the first resource group.
20. The method of claim 18, wherein a first priority of the pre-defined priority is placed on maximizing the net number of outgoing weak affinities satisfied, a second priority is placed on minimizing the number of resource groups removed from one of the plurality of candidate nodes, a third priority is placed on maximizing the net number of incoming weak affinities satisfied, and a fourth priority is placed on the ordinal position of the affinity candidate node in the nodelist property.
21. A computer system for activating resource groups, comprising:
 - a processor;
 - a memory;
 - a storage device; and
 - software instructions stored in the memory for enabling the computer system under control of the processor to:
 - determine a plurality of candidate nodes for the first resource group;
 - eliminate any of the plurality of candidate nodes that violates a strong affinity of the first resource group to generate a plurality of affinity candidate nodes;

determine a composite value for each of the plurality of affinity candidate nodes; and
sort the plurality of affinity candidate nodes based on a pre-defined priority using the composite values to obtain a sorted list.

22. The system of claim 21, further comprising:

software instructions stored in the memory for enabling the computer system under control of the processor to activate the first resource group on at least one of the plurality of affinity candidate nodes based on the sorted list.

23. The system of claim 21, wherein software instruction to determine the composite value comprises:

determining a number of resource groups that are removed from one of the plurality of affinity candidate nodes when the first resource group is activated on the one of the plurality of candidate nodes;
calculating the net number of outgoing weak affinities satisfied when the first resource group is activated on the one of the plurality of candidate nodes;
calculating the net number of incoming weak affinities satisfied when the first resource group is activated on the one of the plurality of candidate nodes;
and
determining an ordinal position of the one of the plurality of candidate nodes in the first resource group.

24. The system of claim 20, wherein a first priority of the pre-defined priority is placed on maximizing the net number of outgoing weak affinities satisfied, a second priority is placed on minimizing the number of resource groups removed from one of the plurality of candidate nodes, a third priority is placed on maximizing the net number of incoming weak affinities satisfied, and a fourth priority is placed on the ordinal position of the affinity candidate node in the nodelist property.